

Insider Threat Control: Understanding Data Loss Prevention (DLP) and Detection by Correlating Events from Multiple Sources

George J. Silowash
Christopher King

January 2013

TECHNICAL NOTE
CMU/SEI-2013-TN-002

CERT[®] Program

<http://www.sei.cmu.edu>



Copyright 2012 Carnegie Mellon University

This material is based upon work funded and supported by Department of Homeland Security under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center sponsored by the United States Department of Defense.

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of Department of Homeland Security or the United States Department of Defense.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

This material has been approved for public release and unlimited distribution except as restricted below.

Internal use:* Permission to reproduce this material and to prepare derivative works from this material for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use:* This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

* These restrictions do not apply to U.S. government entities.

Carnegie Mellon® and CERT® are registered in the U.S. Patent and Trademark Office by Carnegie Mellon University.

DM-0000083

Table of Contents

Acknowledgments	vii
Abstract	ix
1 Introduction	1
1.1 Audience and Structure of this Report	1
2 Mitigating Insider Threat: Tools and Techniques	2
2.1 The CERT Insider Threat Database	3
2.2 The Windows Registry	3
2.3 Controlling USB Devices	4
2.4 Auditing USB Device Usage	4
2.4.1 Create an Auditing Policy	5
3 Identifying Sensitive Data	7
3.1 OpenDLP	7
3.1.1 Requirements	7
3.1.2 Background	7
3.1.3 OpenDLP and Regular Expressions	8
3.1.4 Create a Scan Profile	9
3.1.5 Create a Scan Profile	10
3.1.6 Start a Scan	10
3.1.7 View the Scan Results	11
4 Correlating Audit Events Across Tools, Machines, and Users	14
5 Conclusion	16
6 References	17

List of Figures

Figure 1:	Windows Registry View of the <code>USBSTOR</code> Key	3
Figure 2:	USBDeview USB Device Information	4
Figure 3:	The Process of Creating a New OpenDLP Scan Profile	9
Figure 4:	The Process of Creating an OpenDLP Scan	11
Figure 5:	OpenDLP Scan Status	12
Figure 6:	Results from the Enterprise File Servers Scan	12
Figure 7:	Scan Results for a Particular Machine	13

List of Tables

Table 1:	Instructions for Creating a New Auditing Policy	6
Table 2:	Sample Methods for Identifying Sensitive Information	8
Table 3:	Instructions for Creating a New Scanning Profile	10
Table 4:	Instructions for Creating an OpenDLP Scan	11

Acknowledgments

Special thanks to our sponsors at the U.S. Department of Homeland Security, National Cyber Security Division, Federal Network Security branch for supporting this work.

Abstract

Removable media, such as universal serial bus (USB) flash drives, present unique problems to the enterprise since insiders can use such media to remove proprietary information from company systems. Insiders may do this for legitimate reasons, such as to work on material at home, or they may do so for malicious reasons, such as to steal intellectual property.

Organizations must establish and implement effective methods and processes to prevent unauthorized use of removable media while still allowing users with a genuine business need to access and remove such media. In addition, organizations should establish sound methods to track critical electronic assets so that they may better protect them.

This report focuses on the theft of intellectual property using removable media, in particular, USB devices. We present methods to control removable media devices in a Microsoft Windows environment using Group Policy within an Active Directory environment. We also explore OpenDLP, an open source tool for identifying where sensitive data resides on organizational systems.

1 Introduction

Removable media, such as universal serial bus (USB) flash drives, present unique problems to the enterprise since insiders can use such media to remove proprietary information from company systems. Insiders may do this for legitimate reasons, such as to work on material at home, or they may do so for malicious reasons, such as to steal intellectual property.

The staff members of the CERT[®] Program, part of Carnegie Mellon University's Software Engineering Institute, have seen instances where removable media played a role in a malicious insider's attack. In light of this, organizations must establish and implement effective methods and processes to prevent unauthorized use of removable media while still allowing users with a genuine business need to access and remove such media. In addition, organizations should establish sound methods to track critical electronic assets so that they may better protect them.

This report presents methods to control removable media devices in a Microsoft Windows environment using Group Policy within an Active Directory environment. The report also explores an open source tool, OpenDLP, for identifying where sensitive data resides on organizational systems.

1.1 Audience and Structure of this Report

This report is a hands-on guide for system administrators who are implementing USB device auditing and want to have a better understanding of where sensitive organizational data lives.

This remainder of this technical note is organized as follows:

- Section 2 describes some of the techniques available for establishing proper audit policies and technical controls to reduce the risk of malicious insider activity.
- Section 3 outlines how system administrators can use data loss prevention (DLP) products to help identify the organization's sensitive data.
- Section 4 describes how administrators can use a centralized logging system to correlate audit events across machines, tools, and users.
- Section 5 summarizes this technical note.

[®] CERT is a registered trademark owned by Carnegie Mellon University.

2 Mitigating Insider Threat: Tools and Techniques

Malicious insiders are able to act within an organization by taking advantage of weaknesses they find in systems. Organizations must be aware of such weaknesses and how an insider may exploit them; organizations must also be aware of the many ways in which weaknesses are introduced. For example, an organization may have insecure configurations or have relaxed or nonexistent security policies. In other cases, a lack of situational awareness introduces weaknesses that malicious insiders can exploit. Additionally, an organization that allows its employees to use USB devices is essentially increasing the organization's potential for data leakage. Establishing proper audit policies and technical controls, as discussed in this section, will mitigate some of the risks.

We define a *malicious insider* as a current or former employee, contractor, or business partner who

- has or had authorized access to an organization's network, system, or data
- intentionally exceeded or misused that access
- negatively affected the confidentiality, integrity, or availability of the organization's information or information systems

Our research has revealed that most malicious insider crimes fit under one of three categories: IT sabotage, theft of intellectual property, and fraud. Additionally, a 2011 SEI report titled *Insider Threat Control: Using Centralized Logging to Detect Data Exfiltration Near Insider Termination* presents an example of an insider threat pattern based on the insight that "many insiders who stole their organization's intellectual property stole at least some of it within 30 days of their termination" [1].

This report focuses on the theft of information using removable media, in particular, USB devices. When USB devices are introduced to a Microsoft Windows-based system, the system generates many remnants that can be audited or possibly used for forensic analysis. Therefore, it is important to understand how USB devices interact with the system.

In this section, we present tools and techniques that an organization can implement to mitigate insider threats. We describe the CERT insider threat database and the Windows Registry, outline techniques for controlling USB devices, and present methods for monitoring and auditing USB device usage. The tools and techniques presented in this report represent just a subset of various practices an organization could implement to mitigate insider threats. For example, DLP tools, such as OpenDLP, can scan databases for sensitive information; however, any commercial DLP tool could also complete this activity. Once sensitive information has been identified, information security teams can implement additional security accordingly.

Please note that since OpenDLP is only capable of searching for regular expressions found in cleartext, encryption defeats this tool. Since encryption converts plaintext into an unreadable form, regular expression scanning is rendered useless. Nonetheless OpenDLP is an example of a

simplified DLP tool that has a subset of the capabilities of a COTS tool set. We discuss OpenDLP further in Section 3.1 of this report.

2.1 The CERT Insider Threat Database

The CERT insider threat research is based on an extensive set of insider threat cases that are available from public sources, court documents, and interviews with law enforcement and/or convicted insiders, where possible. The database contains more than 700 cases of actual malicious insider crimes. Each case is entered into the database in a consistent, repeatable manner that allows us to run queries to search for specific information. The database breaks down the complex act of the crime into hundreds of descriptors, which can be further queried to provide statistical validation of our hypotheses. Since the database has captured very granular information about insider threat cases, it provides a way to find patterns of insider activity, discover possible precursors to insider attacks, and discover technical and nontechnical indicators of insider crime. This helps us to establish trends and commonalities and identify techniques that may be helpful in mitigating insider threats.

2.2 The Windows Registry

The Microsoft Windows Registry records a wealth of information when a device is connected to the system. However, to implement in-depth auditing of USB device events, we must first understand what Windows records in the registry. The information shown in Figure 1 was derived from a machine running Microsoft Windows 7 as the operating system. USB device activity is stored in the Registry Key and subkeys.¹

To view USB device information, open the registry editor and navigate to this key:

Computer\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Enum\USBSTOR.
This registry key will be used later in this report to implement an audit policy for USB devices.

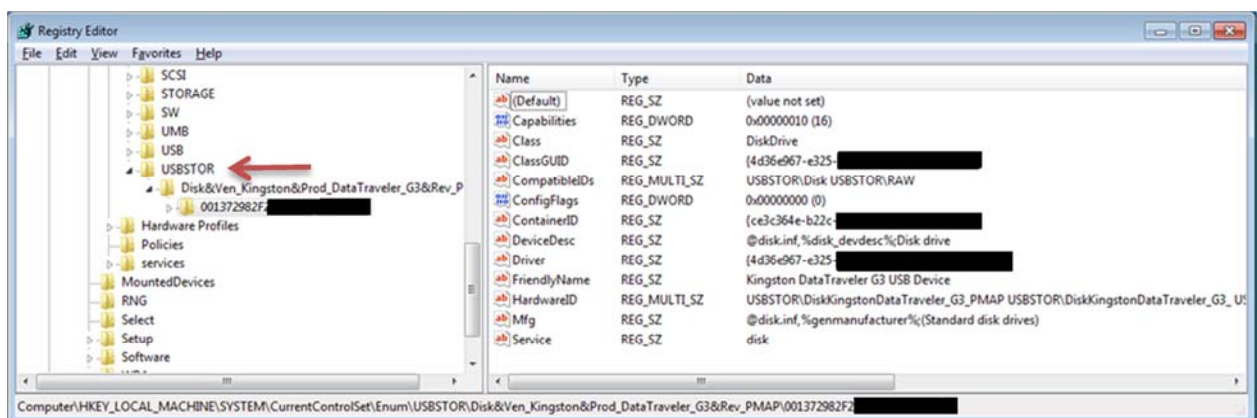


Figure 1: Windows Registry View of the USBSTOR Key


In Figure 1, each subkey under the USBSTOR key was created when a corresponding device was first introduced to the system. When you expand the subkeys beneath USBSTOR, the devices that

¹ Serial numbers and other sensitive information have been redacted from the figures in this document.

have been connected to the system are listed. Also in Figure 1, the USBSTOR key is listed with a subkey of Disk&Ven_Kingston&Prod_DataTraveler_G3&Rev_P.... This key contains the name of the device vendor (Kingston) and product (DataTraveler G3); expanding the key will list additional keys that correlate to the serial numbers (if available) of the devices that are of the same vendor and product type. For more details, see the article titled “USB History Viewing” on the Forensics Wiki [2].

Navigating the Windows Registry to gather related information can be daunting and time consuming. Freeware tools, such as NirSoft’s USBDeview, are available to assist users in accomplishing tasks associated with the Windows Registry [3].

Figure 2 illustrates much of the same information found in Figure 1, but in a format that is easier to read. In addition, by observing the “Created Date” column as well as the “Last Plug/Unplug” event, you can determine when the device was first introduced to the system.



Device Name	Description	Device Type	Connected	Safe To Unpl...	Disabled	USB Hub	Drive Letter	Serial Num...	Created Date	Last Plug/Unplug Da...
0002.0000.0000.001.00...	USB Input Device	HID (Human Interface D...	No	Yes	No	No			10/13/2011 1:18:28 PM	10/23/2011 9:26:45 PM
0002.0000.0000.001.00...	USB Input Device	HID (Human Interface D...	No	Yes	No	No			10/13/2011 1:18:28 PM	10/23/2011 9:26:45 PM
Port_#0001.Hub_#0001	USB Composite Device	Unknown	Yes	Yes	No	No			10/13/2011 1:18:28 PM	10/23/2011 9:26:44 PM
Port_#0001.Hub_#0002	Kingston DataTraveler G3 USB Device	Mass Storage	No	Yes	No	No	E:	001372982F...	10/14/2011 12:19:57 PM	10/18/2011 7:27:57 PM

Figure 2: USBDeview USB Device Information

2.3 Controlling USB Devices

In a Microsoft Windows Server 2008 and Windows Vista (or higher) environment, you can manage USB devices using Group Policy Objects (GPOs) and Active Directory. This is useful for organizations that want to block or control specific actions that a user can perform with a USB device. If USB devices will be used within the organization, CERT staff recommend limiting access to approved devices owned by the organization.

For a more in-depth discussion on this topic, see the Microsoft article titled “Step-By-Step Guide to Controlling Device Installation Using Group Policy” [4].

2.4 Auditing USB Device Usage

You can use Microsoft Windows Server 2008 and Windows 7 clients to configure an auditing policy for removable devices, and you can use Group Policy to apply audit settings to systems within an organization.² USB device activities are just one of the many events that you should audit; you should also configure systems to audit other events, such as sensitive file accesses or changes, user account activity, and changes to system configuration or policy.

If USB devices are permitted to be used within the organization, it may be difficult to determine which events are high priority. Therefore, an organization must identify sensitive or high-risk data

² Other operating systems may be supported; however, these were the only systems tested in the CERT insider threat lab.

and then design audit and alerting systems that correlate the access of such information to the usage of USB devices. For example, if file-access auditing is configured to flag sensitive files in tandem with using USB device auditing, it would be possible to design a SIEM rule that provides an alert when sensitive data is accessed around the time of USB device usage.

2.4.1 Create an Auditing Policy

To create an audit policy, open the Group Policy Management tool on a server or administrator workstation using an administrative account. To do this, select Start > Run and type the following:

```
mmc gpmmc.msc
```

This will open the Group Policy Management Console. The remaining steps are outlined in Table 1; these steps enable auditing across all domain computers from which the policy is linked. Once Group Policy refreshes across the machines in the organization, you will begin to see audit events appear in the security logs of the client machines when a USB device is used. The user can then forward the USB device audit events to a Security Information and Event Management (SIEM) system or other central logging server for analysis and correlation. Microsoft Windows Server 2008 and Windows Vista and higher support the forwarding of events to a centralized server natively. Some SIEM systems may require you to install a software client on the system to collect and forward events to the SIEM system.

For a deeper introduction to this topic, please see the Microsoft article titled “Windows Event Collector” [5].

Table 1: Instructions for Creating a New Auditing Policy

1. Expand the Group Policy Objects (GPO) folder.
2. Right-click on the GPO folder and click the "New" key.
3. Enter a name for the GPO (e.g., USB Device Audit Settings).
4. Click the "OK" key.
5. Expand the following Policy Object:
Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy
6. On the right side of the Management Console, double-click the "Audit Object Access" policy option.
7. Add a checkmark to each of the following options:
 - a. Define these policy settings
 - b. Success
 - c. Failure
8. Click the "OK" key.
9. Expand the following registry key:
Computer Configuration\Windows Settings\Security Settings\Registry
10. Right-click on the right side of the Management Console screen and select "Add Key."
11. Navigate to the following key in the "Select Registry Key" dialog box:
MACHINE\SYSTEM\CurrentControlSet\Enum\USBSTOR
12. Click the "OK" key. (Doing so will display the Database Security window.)
13. In the Database Security window, click the "Advanced" button.
14. Click the "Auditing" tab.
15. Click the "Add" key.
16. In the "Enter object name to select" text box, type Everyone.
17. Press the "Enter" key.
18. In the "Full Control" area, check the following boxes:
 - a. Successful
 - b. Failed
19. Click the "OK" key four times.
20. Close the Group Policy Management Editor window. (Doing so will display the Group Policy Management console.)
21. Link the new GPO object to the domain.

3 Identifying Sensitive Data

Many commercial DLP products can identify sensitive data within an organization. Data loss prevention products allow organizations to establish policies for how data should be protected in the following circumstances:

- at rest (e.g., data stored on a hard drive)
- in motion (e.g., data traveling on a network)
- in use (e.g., data being used by someone who is accessing and modifying files)

This information can then be used to mitigate risks associated with data exfiltration.

In the course of our research, we evaluated OpenDLP, an open source DLP product to determine if it has merit in mitigating insider threat. We discuss OpenDLP in the next section.

3.1 OpenDLP

OpenDLP, an open source data loss prevention tool, can scan databases for sensitive information. As mentioned earlier, since OpenDLP is only capable of searching for regular expressions found in cleartext, encryption defeats this tool. Since encryption converts plaintext into an unreadable form, regular expression scanning is rendered useless. Nonetheless OpenDLP is an example of a simplified DLP tool that has a subset of the capabilities of a COTS tool set.

3.1.1 Requirements

We assume that you are familiar with and have downloaded and installed VirtualBox, Oracle's open source, virtual machine software.³ VirtualBox is required in order to use the pre-built OpenDLP virtual machine. In addition, you will need to download the OpenDLP virtual machine image files to a single directory and unzip them.⁴ Source code is also available to download and install should you choose to use an existing Linux-based system or if you prefer not use virtual machines. The virtual machine option allows you to quickly configure a system with little additional configuration; installing from source code requires a large amount of configuration and other software dependencies. In addition, the `README-VM.txt` file that is bundled with the virtual machine images contains instructions that you must follow to configure the software properly.⁵

3.1.2 Background

OpenDLP is an open source software package that assists organizations in identifying sensitive information that is left unsecured at rest. OpenDLP can scan devices using both an agent and

³ Please note that we tested the process of importing the OpenDLP VirtualBox-based machine into a VMware Workstation, but it failed to import correctly.

⁴ You can find the virtual machine image files online (<https://code.google.com/p/opendlp/>) [7].

⁵ While outside the scope of this report, we encourage users to change all default passwords and certificates to avoid security risks.

agentless approach. (Agent and agentless scans are discussed further in Section 3.1.4.1.) The user can configure the tool to search documents for specific phrases that may identify sensitive information within the organization. The user can also create different scanning profiles to search for specific phrases.

OpenDLP, while not an active defender of sensitive organizational data, does identify basic Microsoft Office 2010 documents and other zip files containing sensitive information. The reports generated from this tool can be used to further mitigate insider threat through the use of access control lists (ACLs) and auditing.

3.1.3 OpenDLP and Regular Expressions

OpenDLP contains built-in regular expressions (RegExs) for all major credit cards and social security numbers.⁶ The feedback that we hear from practitioners is that DLP tools work on well-formatted data such as social security numbers and credit card numbers; however, DLP approaches are not as reliable when you are working with proprietary information, intellectual property, or other sensitive data. It is therefore recommended that companies review their intellectual property and other data to devise a list of keywords that may flag a document as sensitive. Table 2 lists examples of such keywords; also included in this table are associated regular expressions that should be added to the OpenDLP RegEx library. Any number of regular expressions can be crafted for a specific need within the organization. The implementer only needs to have a basic understanding of regular expressions to create new expressions.⁷

To add a regular expression to the library, click the Regular Expressions menu on the left side of the screen, and then click Create New RegEx and provide a name. The name should contain only letters and numbers. Hyphens (-) and underscores (_) are also permitted; however, spaces are not permitted.

Table 2: Sample Methods for Identifying Sensitive Information

Keywords	Regular Expression
Company Confidential	(?ism) COMPANY\sCONFIDENTIAL
Company Proprietary	(?ism) COMPANY\sPROPRIETARY
Confidential	(?ism) CONFIDENTIAL
Controlled Unclassified Information	(?ism) CONTROLLED\sUNCLASSIFIED\sINFORMATION
CUI (acronym for <i>Controlled Unclassified Information</i>)	(?ism) CUI
For Official Use Only	(?ism) FOR\sOFFICIAL\sUSE\sONLY
FOUO (acronym for <i>For Official Use Only</i>)	(?ism) FOUO
Limited Distribution	(?ism) LIMITED\sDISTRIBUTION
<i>Project Name</i>	(?ism) PROJECT\sNAME
Proprietary	(?ism) PROPRIETARY
Secret	(?ism) SECRET
Top Secret	(?ism) TOP\sSECRET

⁶ Visit the opendlp wiki to read more about the Regular Expressions that are used in OpenDLP (<http://code.google.com/p/opendlp/wiki/RegularExpressions>) [7].

⁷ The regular expressions listed in Table 2 were designed to detect keywords within the headers and footers of sensitive documents. They may produce additional results depending on the context of the document.

3.1.4 Create a Scan Profile

Once you have created the regular expressions that will be used to scan for sensitive information, you should create a scan profile. To do this, click the “Profiles” option on the left side of the screen and then select “Create New Profile.” Figure 3 illustrates this process.

Create a new scan profile	
Profile Name	SMB-Test
Scan Type	<div>Windows Filesystem (agent) Windows Filesystem (agentless over SMB) Windows Network Share (agentless over SMB) UNIX Filesystem (agentless over SSH) Microsoft SQL Server (agentless) MySQL (agentless)</div>
Mask Sensitive Data?	
Username	
Password	

Figure 3: The Process of Creating a New OpenDLP Scan Profile

3.1.4.1 Agent vs. Agentless Scanning

OpenDLP can perform agent or agentless scanning. An agent-based scan deploys a software package to a client that searches for sensitive data and returns the results to the OpenDLP server. Agentless scans are conducted by the OpenDLP server, where the results are processed and stored.

In the CERT insider threat lab, we have been more successful using agentless server message block (SMB) scans using small numbers of Internet protocol (IP) addresses at one time. This scenario is similar to what you will likely find in most organizations. For example, organizations typically require users to store information on servers. Often times these servers number less than the workstations and therefore could be one of the targets of an agentless SMB scan. Because of these considerations, the remainder of this report will focus on agentless SMB scans.

SMB scans require an administrative account and access to default or administrative shares (C\$, D\$, etc.) on the target systems. Typically, by default, the administrative shares are enabled with full control given to members of the “Domain Administrators” group.⁸ However, to enhance security, some organizations have disabled these administrative shares.

If administrative shares are not available, a network share scan can be configured in much the same way as an administrative SMB scan. Modifications to the processes below will be noted for network share scans. In addition, Windows firewall configurations may block these scans. Therefore, you may need to adjust the firewall policy on the targeted systems to permit network traffic

⁸ For further discussion of administrative shares, please visit the “Disable Administrative Shares” web page on the Petri website (http://www.petri.co.il/disable_administrative_shares.htm) [7].

from the OpenDLP scanner. We suggest that you conduct scans during non-production hours when possible. We also suggest that you scan only a small number of machines at a time.

3.1.5 Create a Scan Profile

A scan profile creates a template from which a scan will be started later. The profile contains Windows account information and the regular expressions that will be included in searches.

To create a new scan profile, follow the steps outlined in Table 3. (If a specific option is not discussed in the table, assume the default settings listed.)

Table 3: Instructions for Creating a New Scanning Profile

- | |
|--|
| <ol style="list-style-type: none">1. Enter a profile name in the "Profile Name" field.2. Set the scan type to "Windows Filesystem (agentless over SMB)." (For a network share scan, select "Windows Network Share...")3. Select whether or not you want to mask the sensitive data in the scan results.4. Enter the administrator's username for the targets of the scan. This may be a local administrator or domain administrator. Note that if you are scanning multiple systems, the username and password must be the same on all systems. Therefore, in a domain environment, we advise you to use a domain administrator account.5. Enter the password for the account.6. Enter the domain or workgroup name for the above account.7. Indicate which client directories to scan. For a Windows Share scan, list the directories within the share you want to scan (e.g., <code>users\home\</code>). You may also click the question mark icon next to the directories to scan the input box for additional information.8. Set the File Extensions option to "Scan all files."<ul style="list-style-type: none">– Selecting this option prevents files from going undetected. To speed up the scan, you may wish to customize the exception list. However, simply because a file is saved with a particular extension does not necessarily mean that file contains that type of data. (For example, to avoid simplistic detection systems, a Microsoft Word document could be saved with a ".JPG" extension.)9. Select the regular expressions to include in the scan.<ul style="list-style-type: none">– While some regular expressions may not apply to a particular system, you may want to include them anyway to flag possible policy violations or data spillage. For example, to detect data spillage on an unclassified system, you may want to include RegExs for Top Secret, Secret, and Confidential.10. Set the concurrent deployments to a number between 1 and 100.<ul style="list-style-type: none">– You will need to adjust this number accordingly depending on 1) your environment, 2) the system resources available to the OpenDLP virtual machine, and 3) the options selected in the scan. We advise you to start with a small number and gradually increase it.11. Click the "Submit" button to save the scanning profile. |
|--|

3.1.6 Start a Scan

Once you have built a scanning profile, you can then start a scan. To do this, follow the steps outlined in Table 4, and illustrated in Figure 4.

Table 4: Instructions for Creating an OpenDLP Scan

1. Enter a scan name in the “Scan Name” field.
2. Select the appropriate scanning profile.
3. Enter the IP addresses to include in the scan.
4. Click the Start button.

If you have configured the scanning profile correctly and entered valid IP addresses, the scan should start. Immediate logging information related to scan success or failure will be displayed on this screen. Detailed logging is available in the scan results.

OpenDLP 0.4.2

- Main
- Profiles
- Regular Expressions
- Scans
 - Start New Scan**
 - View Scans/Results
 - Export Scan Results
 - Delete Scan Results
- False Positives
- Logs
- OpenDLP Homepage

Start a New Scan

Scan name	Enterprise_File_Servers
Profile	SMB-User ▼ (or create a new profile)
Notes	If you are doing a Windows Share scan, enter systems in this format: \\1.2.3.4\Share Otherwise, just list IP addresses.
Systems to scan (newline-delimited)	<div style="border: 1px solid #ccc; padding: 5px; min-height: 100px;"> 192.168.1.3 192.168.1.4 192.168.1.5 </div>
<input type="button" value="Start"/>	

Figure 4: The Process of Creating an OpenDLP Scan

3.1.7 View the Scan Results

Once the scan has completed, the results will be available on the “View Results” page. You can reach this page by clicking on the “Scans” menu option and then selecting “View Scans/Results.” This process is illustrated in Figure 5.

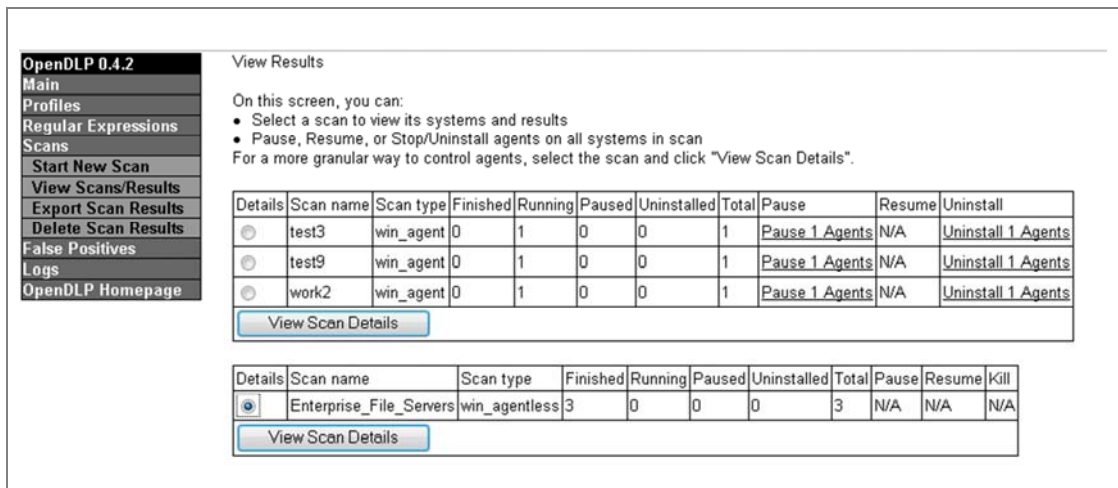


Figure 5: OpenDLP Scan Status

To view details of sensitive information found on a particular machine, select any of the IP addresses. In the example shown in Figure 6, the machine associated with IP address 192.168.1.3 returned 77 findings.

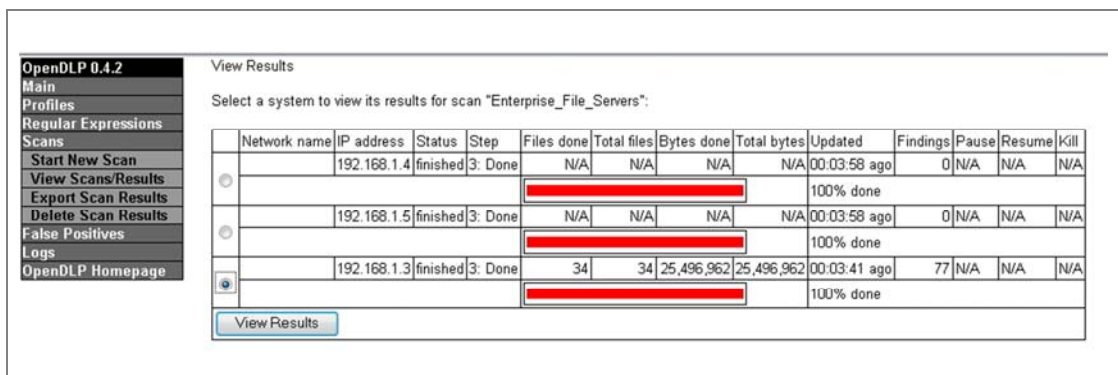


Figure 6: Results from the Enterprise File Servers Scan

Figure 7 depicts an excerpt from the scan results, where you will see a number of documents that contain sensitive information. If you click a filename, you will be given the opportunity to open the file for further examination.

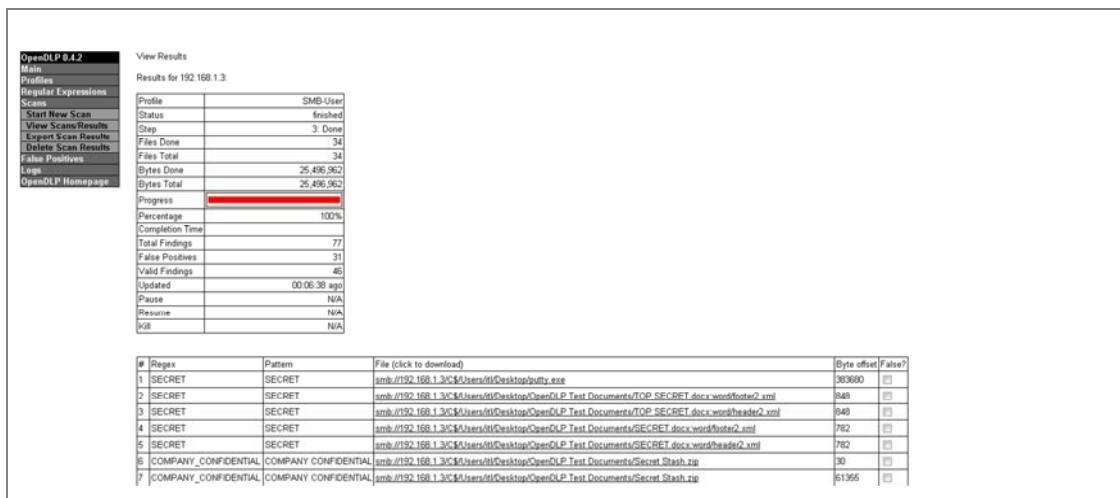


Figure 7: Scan Results for a Particular Machine

Once you have verified that the files identified contain sensitive information, you should take further investigative steps to determine who has had access to the files. You should also determine if any of the ACLs for the files need to be changed to prevent unauthorized disclosure, modification, or deletion. The ACLs should be modified so that only authorized users with a genuine business need or need to know have access to them. Take care to review group membership and permissions; failure to do so may create a backdoor for an unauthorized user to access a file.

The following example illustrates a backdoor related to access rights.

Joe Smith is a member of an organization's Sales Security group; members of this group are allowed to access the master customer file. However, as part of his daily job, Joe does not need to have access to this file. Because of this, the organization's information security team has not listed him in the ACL for this file; however, since he is a member of the Sales Security group, he does, in fact, have access. This creates a backdoor that allows Joe unnecessary access to the master customer file.

4 Correlating Audit Events Across Tools, Machines, and Users

By implementing a centralized logging system, audit events can be correlated across tools, machines, and users. Taking this approach allows you to use a SIEM system to quickly identify related events that may be of concern. A SIEM system rapidly processes numerous events and is capable of generating alerts when it detects patterns of suspicious behavior.

For a more complete picture of an organization's data, a DLP tool should provide continuous monitoring of sensitive information and forward security events to the SIEM system for review. The DLP tool should address data at rest, in motion, and in use.

Although OpenDLP only addresses data at rest, we used it in this report to raise awareness of where sensitive data lives within organizations. You can use the information from Section 2.2, The Windows Registry, to search for specific devices in log files or on a SIEM system.

A SIEM solution should be capable of generating an alert based on events surrounding the use of a USB device. Examples of the types of information the SIEM system can capture include user account information, USB device audit events, and alerts generated by a DLP product. Consider the following actual case:

An insider was employed as a senior financial analyst by a victim organization, a financial institution. Every Sunday, the insider entered the organization's offices and downloaded 20,000 mortgage applicant records to a USB flash drive. The insider sometimes downloaded the records during normal working hours.

Members of the organization noticed that the insider had been coming to work outside of normal working hours, but they believed the insider was merely a hard-working employee.

The organization has a policy that prohibits flash drives or other storage devices from being used on its computers. The organization believed it had disabled flash drive access on all computers; however, the insider located and used the organization's one flash-enabled computer.

Over a two-year period, the insider downloaded and sold over 2,000,000 records that contained personally identifiable information (PII). As a result, at least 19,000 mortgage applicants became victims of identity theft, and dozens of class action lawsuits have been filed against the victim organization.

In this scenario, USB audit logs generated by modifications to the Windows Registry, Windows File Access Auditing, and a DLP tool would have revealed that someone mounted a USB device and accessed and copied sensitive information from the system. This would have prompted the organization's information security team to conduct further analysis of the events in question to determine whether an actionable security incident took place. Further analysis would have revealed that a USB device had been mounted; shortly thereafter, the DLP would have generated a

log and sent it to the SIEM system. Events related to PII data should have a very high priority rating as this data is very sensitive and often protected by laws and regulations. The log would have shown that data had moved from a server to a client and then onto the removable media. The SIEM system would have used a rule to detect activities such as USB insertion and data that moved from a server to a desktop and then to a USB device (or simply from a server to a USB device). This rule could have triggered an alert for the organization's information security team, which would have conducted further analysis.

In the above scenario, had the financial institution implemented proper insider threat mitigation strategies, the insider may have been detected much earlier or stopped completely.

5 Conclusion

DLP solutions assist organizations in discovering sensitive company information, and this allows information security teams to monitor data usage and detect and mitigate insider threats.

DLP tools also afford data owners a deeper understanding of which access control lists need to be in place to determine who has access to sensitive files and what their security permissions allow them to do with the files. The use of removable media within an organization must be carefully considered and included as part of the organization's risk assessment process. By evaluating what users can do with USB devices within the organization and what files can be accessed and copied to USB removable media, organizations may be more inclined to increase security controls that prohibit, restrict, and monitor USB device activity.

Organizations can use the information generated by implementing the USB device auditing policy described in Section 2.4 of this report. Information collected from OpenDLP scans across the enterprise will help to identify risks and help the organization understand whether it needs to implement additional mitigation strategies.

6 References

- [1] M. Hanley and J. Montelibano. "Insider Threat Control: Using Centralized Logging to Detect Data Exfiltration Near Insider Termination," Software Engineering Institute, Carnegie Mellon University, Pittsburgh, 2011.
<http://www.sei.cmu.edu/library/abstracts/reports/11tn024.cfm>
- [2] Forensics Wiki. "USB History Viewing," October 2011.
http://www.forensicswiki.org/wiki/USB_History_Viewing
- [3] NirSoft. "NirSoft Website." http://www.nirsoft.net/utils/usb_devices_view.html
- [4] D. Bishop, "Step-By-Step Guide to Controlling Device Installation Using Group Policy," June 2007. <http://msdn.microsoft.com/en-us/library/bb530324.aspx>
- [5] Microsoft Corporation. "Windows Event Collector." [http://msdn.microsoft.com/en-us/library/windows/desktop/bb427443\(-vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/bb427443(-vs.85).aspx)
- [6] Petri IT Knowledgebase. "Disable Administrative Shares."
http://www.petri.co.il/disable_administrative_shares.htm
- [7] OpenDLP. "OpenDLP, Data Loss Prevention Suite." <http://code.google.com/p/opendlp/>

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE January 2013		3. REPORT TYPE AND DATES COVERED Final
4. TITLE AND SUBTITLE Insider Threat Control: Understanding Data Loss Prevention (DLP) and Detection by Correlating Events from Multiple Sources			5. FUNDING NUMBERS FA8721-05-C-0003	
6. AUTHOR(S) George J. Silowash & Christopher King				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213			8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2013-TN-002	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) ESC/CAA 20 Schilling Circle, Building 1305, 3rd Floor Hanscom AFB, MA 01731-2125			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS			12B DISTRIBUTION CODE	
13. ABSTRACT (MAXIMUM 200 WORDS) Removable media, such as universal serial bus (USB) flash drives, present unique problems to the enterprise since insiders can use such media to remove proprietary information from company systems. Insiders may do this for legitimate reasons, such as to work on material at home, or they may do so for malicious reasons, such as to steal intellectual property. Organizations must establish and implement effective methods and processes to prevent unauthorized use of removable media while still allowing users with a genuine business need to access and remove such media. In addition, organizations should establish sound methods to track critical electronic assets so that they may better protect them. This report focuses on the theft of intellectual property using removable media, in particular, USB devices. We present methods to control removable media devices in a Microsoft Windows environment using Group Policy within an Active Directory environment. We also explore OpenDLP, an open source tool for identifying where sensitive data resides on organizational systems.				
14. SUBJECT TERMS Data loss prevention, DLP, insider threat, removable media, universal serial bus, USB, malicious insiders			15. NUMBER OF PAGES 29	
16. PRICE CODE				
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	